

CLAIMS

What is claimed is:

1 A method for joining an inner table and an outer table in a database in response to a  
2 query statement having an inclusion operator, said method comprising:

3 (a) providing an inner table;

4 (b) providing an outer table having zero or more records for inclusion with said  
5 inner table;

6 (c) transposing said outer table and said inner table to form a right table and a  
7 left table, respectively;

8 (d) creating a hash table from said left table;

9 (e) obtaining a hash value from said right table;

10 (f) probing said hash table with a said hash value from said right table to  
11 determine if said hash value matches a value in said hash table; and

12 (g) if said outer hash value matches said value in said hash table, then evaluating  
13 the actual values and, if the join condition is satisfied, including said row of said outer table  
14 in a result.

1 2. The method of claim 1, wherein scanning said inner table is accomplished one row  
2 at a time per processor.

1 3. The method of claim 1, wherein said step of obtaining a hash value from said right  
2 table includes calculating a hash value.

1 4. The method of claim 1, wherein said step of obtaining a hash value from said right  
2 table includes retrieving a pre-calculated hash value from a row header.

1 5. A method for joining an inner table and an outer table in a database in response to a  
2 query statement having an exclusion operator, said method comprising:

3 (a) providing an inner table;

4 (b) providing an outer table having zero or more records for inclusion with said  
5 inner table;

6 (c) transposing said outer table and said inner table to form a right table and a  
7 left table, respectively;

8 (d) creating a hash table from said left table;

9 (e) obtaining a hash value from said right table;

10 (f) probing said hash table with a said hash value from said right table to  
11 determine if said hash value matches a value in said hash table;

12 (g) if said right hash value do not match any value in said hash table, then  
13 including said row from said right table in a result.

1 6. The method of claim 5, wherein scanning said outer table is accomplished one row  
2 at a time per processor.

1 7. The method of claim 5, wherein said step of obtaining a hash value from said right  
2 table includes calculating a hash value.

FIG. 10 - TABLE 10

1 8. The method of claim 5, wherein said step of obtaining a hash value from said right  
2 table includes retrieving a pre-calculated hash value from a row header.

1 9. A database system having a database with at least two tables for responding to SQL  
2 statements having an inclusion join condition that designate an inner table and an outer table  
3 within said database, said database system comprising:

4 system memory, said system memory constructed and arranged to store a hash table;

5 a processor operative with said system memory, said processor constructed and  
6 arranged to generate a hash table from said inner table, said processor further constructed  
7 and arranged to scan at least one row of said outer table, obtain an outer hash value, and  
8 compare said outer hash value to a hash value in said hash table to determine if said values  
9 match and, if so, evaluating actual values from said rows to determine if a join condition is  
10 satisfied and, if so, then adding said row from said outer table to a result.

1 10. A database system having a database with at least two tables for responding to SQL  
2 statements having an exclusion join condition that designate an inner table and an outer table  
3 within said database, said database system comprising:

4 system memory, said system memory constructed and arranged to store a hash table;

5 a processor operative with said system memory, said processor constructed and  
6 arranged to generate a hash table from said inner table, said processor further constructed  
7 and arranged to scan at least one row of said outer table, obtain an outer hash value, and  
8 compare said outer hash value to a hash value in said hash table to determine if said values  
9 match and, if not, then adding said row from said outer table to a result.

1 11. A database system for accessing a database in response to a SQL statement having  
2 an inclusion condition, said database system comprising:  
3 a massively parallel processing system comprising:  
4 one or more nodes;  
5 a plurality of CPUs, each of the one or more nodes providing access to one  
6 or more CPUs;  
7 a plurality of virtual processes each of the one or more CPUs providing  
8 access to one or more virtual processes;  
9 each virtual process configured to manage data stored in one of a plurality of  
10 data-storage facilities;  
11 a parsing engine comprising:  
12 a session control, which controls access to the plurality of virtual processes;  
13 a parser, which interprets and optimizes a command for access to a database  
14 to produce a series of steps to be performed; and  
15 a dispatcher, which controls the sequence of the series of steps and  
16 determines which of the plurality of virtual processes will perform each of the series  
17 of steps;  
18 where the parser includes:  
19 an optimizer which optimizes a plan for executing the series of steps;  
20 where the optimizer includes:  
21 an inclusion hash join capability; and

22 capability to recognize an inclusion merge join and replace with said  
23 inclusion hash join having an inner table and an outer table, wherein said inner table  
24 and said outer table are assigned as a left table and a right table, respectively, and  
25 said left table is hashed into a hash table and hash values from said right table are  
26 used to probe said hash table to determine if an inclusion condition is satisfied for a  
27 row of said right table and, if so, then evaluating a join condition and, if said  
28 evaluated indicates said inclusion condition is satisfied, then including said row of  
29 said right table in a result.

1 12. A database system for accessing a database in response to a SQL statement having  
2 an exclusion condition with an inner and an outer table, said database system comprising:

3 a massively parallel processing system comprising:

4 one or more nodes;

5 a plurality of CPUs, each of the one or more nodes providing access to one  
6 or more CPUs;

7 a plurality of virtual processes each of the one or more CPUs providing  
8 access to one or more virtual processes;

9 each virtual process configured to manage data stored in one of a plurality of  
10 data-storage facilities;

11 a parsing engine comprising:

12 a session control, which controls access to the plurality of virtual processes;

13 a parser, which interprets and optimizes a command for access to a database  
14 to produce a series of steps to be performed; and

15 a dispatcher, which controls the sequence of the series of steps and  
16 determines which of the plurality of virtual processes will perform each of the series  
17 of steps;

18 where the parser includes:

19 an optimizer which optimizes a plan for executing the series of steps;

20 where the optimizer includes:

21 an exclusion hash join capability; and

22 capability to recognize an exclusion merge join and replace with said  
23 exclusion hash join wherein said inner and outer tables are assigned to a left table  
24 and a right table, respectively, said left table is hashed into a hash table and a row of  
25 said right table is hashed to obtain a hash value that used to probe said hash table to  
26 determine if said hash values match and, if no hash values match, then including said  
27 row of said right table in a result.

407650-1234567890